Fermilab

TECHNICAL SCOPE of WORK

Between

The NOvA Experiment

and

The Fermilab Computing Sector

for

Support of Computing used in the Operation of theNOvA Experiment

3-Feb-2014
Version 1.0

Abstract:

This document is the Technical Scope of Work (TSW), formerly known as a Memorandum of Understanding (MOU), between the Fermilab Computing Sector (CS) and the NOvA experiment for support of the Computing Systems used by the NOvA experiment. This document is intended to clarify the roles and responsibilities of the two parties in supporting the computing resources based upon the requirements agreed to at the time of publication.

Table of Contents

1 Introduction

This document is the Computing Sector – NOvA Technical Scope of Work (TSW) that describes in more detail than the NOvA –Fermilab TSW, the responsibilities of the Fermilab Computing Sector (CS) and the NOvA experiment personnel for computing services used by the experiment. The TSW:

- Will be reviewed on a yearly basis by all the parties to the agreement and amended as requirements change.
- Shall be valid until the end of data analysis for the NOvA experiment.
- Shall cover the long-term computing needs of the experiment including any data preservation needs.
- Shall reflect the computing requirements provided each year in the Computing Sector Strategic and Tactical plans, to which the NOvA experiment provides substantial input.
- Shall refer to the requirements for computing capacity and hardware covered in separate NOvA Computing Requirements documents
- Shall not include activities funded under the NOvA experiment project funds.

The following organizational units are involved in support activities under this TSW:

- The Computing Sector (CS), including the Office of the Chief Information Officer (OCIO), the Core Computing Division (CCD), and the Scientific Computing Division (SCD).
- The NOvA experiment. NOvA offline/production computing groups, NOvA online/data acquisition group, NOvA electronics and electrical engineering support group, NOvA Database support group, the NOvA run coordinator and the NOvA physics analysis groups.

Contacts:

- NOvA: Andrew J. Norman - Computing Sector Liaison to the NOvA experiment, TBA
- Computing Sector: Brian Mckittrick - Service Level Manager, OCIO.

1.1 Overview of Computing Sector Support

Computing Sector service support is provided as specified in the *FNAL Foundation Service Level Agreement (SLA)* [1], which applies to all Computing Sector supported services, except as amended by service-specific Service Level Agreements (SLAs). It is important to note that in general:

- Computing Sector support is provided on an 8x5 basis unless otherwise specified and agreed.
- Additional Service Level Agreements apply for specific services (such as Networking, Database, Grid and Cloud Computing, Storage, Engineering, etc.). These additional SLAs are published in the Service Level Management (subtopic of ITIL Processes and Functions) topic in CS-DocDB.
- All services provided by the Computing Sector are managed through the Computing Sector Service Desk (http://servicedesk.fnal.gov/, or 630-840-2345).

In the event of issues with any service, NOvA experiment personnel shall utilize the Service Desk interface to report any issues. For off hours (outside of the standard 8x5 business hours of Monday-Friday, 8AM to 5PM), the support escalation procedure is to telephone the service desk at 630-840-2345 and select the option to page the on-call service desk personnel.

Computing at Fermilab is governed by the *Fermilab Policy on Computing.*[2] This policy covers all Fermilab-owned computers and any computer, regardless of ownership, when it is connected to the Fermilab network (and/or showing a Fermilab address).

Significant Computing Sector change and maintenance activities shall be coordinated with the experiment so as not to adversely affect operations. Similarly, the experiment shall advise and consult with the Computing Sector prior to performing activities that might result in unusual usage patterns or impose unusually large loads on computing systems.

1.2 Overview of NOvA Experiment Services and Activities

The details of the NOvA systems are documented in the *NuMI Off-Axis nu_e Appearance Experiment Technical Design Report (TDR).*[3] [FERMILAB-DESIGN-2007-01]. Below we summarize the major points to provide a context for the set of services that require operational support.

2 Core Computing Services

2.1 Authentication and Directory Services[4]

The NOvA collaboration will utilize the standard Authentication and Directory Services offerings:

- Kerberos Authentication Services
- Windows Domain Authentication Services
- Services Domain Authentication Services

These services will be provided under the standard Authentication and Directory Services SLA.

In addition the NOvA collaboration has a dedicated Kerberos KDC server, deployed at the Ash River site. This service will have 24x7 support as required for the data acquisition system. (need to reconcile support level with FEF).

2.2 Backup and Restore[5]

The NOvA collaboration will utilize the standard Backup and Restore Services

- Using TIBS backups for NOvA DAQ teststand

2.3 Central Web Hosting[6]

The NOvA collaboration has its main web pages on the standard Central Web Hosting Services and these will be supported.

2.3.1 Apache Central Web Server - Shared Virtual Host

Note this SLA provides 24x7 support for this offering. NoVA has the following websites covered by this SLA

- http://www-nova.fnal.gov/.

The NOvA collaboration has additional web pages which are hosted outside of the central web hosting series including:

- http://novadaq.....(fill in pages)

2.4 Data Center Services[7]

The NOvA collaboration will utilize the standard Data Center Services.

This includes rack space and infrastructure support for:

- 1 rack of DAQ computing in LCC (in support of the NDOS detector)
- 1 rack of DAQ computing in FCC (in support of the Near detector)

2.5 Database Hosting[8]

The NOvA experiment employs databases that are used to store information regarding the operations, calibration and alignment of the detectors. These databases are used by both online and offline systems with different access patterns, replication needs and uptime requirements. Details of all databases used only by the NOvA Project and Experiment is found in NOvA-doc-10602. The relevant information for this documents are characterized in the sections below.

In addition to the database resources located at Fermilab, database servers are installed at the Far Detector site. The requirements for supporting the installation, commissioning, and operation of the Near Detector will be similar to those for the Far Detector. Resources will be provisioned accordingly as Far Detector installation begins.

CS provides tiered levels of support for database services ranging from 8x5 to 24x7, the choice depending upon the application and the type of underlying database. For support under the Database Hosting SLA, it is required that each production database instance will be accompanied by corresponding integration and development instances, both of which receive the lowest available tier of support.

The NOvA collaboration will utilize the standard Database Hosting Services together with the following enhanced services:

2.5.1 Enhanced postgres:

- **NOvA online**

The "online" databases consist of PostgreSQL database servers located on local FNAL sub-networks for each detector Data Acquisition (DAQ) and Detector Controls System (DCS). In the case of the Far detector, the machines are physically located at Ash River. The DCS databases are used to record configuration data that are required to configure/start runs as well as summary information that is written back to the databases to describe the data that was taken. The DCS databases also record configuration data as well as environmental (temperatures, humidity, etc.) data and alarms at the detector location. The DAQ Scientific Linux computing systems are maintained by the SCD FEF department under an SLA for support of the computing systems. The DAQ Microsoft Windows computing systems are maintained by the CCD SOS department under an SLA for support of the computing systems.

- **NOvA offline**

The "offline" database consists of a central PostgreSQL database located at FNAL and is used by the NOvA experiment to store time-dependent calibration, alignment and other quantities needed for offline reconstruction and analysis of data collected in the detectors. The offline database requires 24x7 support, since grid jobs require 24x7 access to these services.

- **Database replication**

The online database tables must be replicated to the offline central database hosted at FNAL, since these data are needed for offline processing of data. Replication should be done in real time, data recorded in an online database table should be replicated within minutes to the offline database.

2.5.2 Responsibilities

2.5.2.1 Computing Sector responsibilities

Install and maintain databases, database servers and applications needed to store and utilize the following mission critical data:

1. Data relating to the construction of the NOvA detector in the FNAL hardware database;
2. Detector configuration and operational data in the DAQ and DCS [online] databases;
3. Offline data, such as beam parameters, detector calibrations, detector alignment, etc.

2.5.2.2 NOvA responsibilities

1. Enter the content of all databases;
2. Interfacing NOvA software with the database applications;
3. Ensuring that users are informed as to appropriate usage patterns, and otherwise assisting CS personnel in investigating and addressing operational issues.
4. For cases in which there is no existing schema or database application, specify and document the requirements, the use cases and queries needed, etc., as requested by the CS.
5. Provide time windows during which regular database maintenance may be performed and security patches applied in a manner consistent with Fermilab security policies and the NOvA Minor Application Plan.

2.5.2.3 Joint responsibilities

Developing and approving the specifications for user access, the database applications and schemas.

Developing and approving the specifications for replication of data between online and offline databases.

Participate in annual "Taking Stock" meetings to long-term operational issues and resource planning. CS will coordinate these meetings.

2.6 Desktop Services[9]

The NOvA collaboration will utilize the standard Desktop Services together with the following enhanced services:

2.7 FermiMail[10]

The NOvA collaboration will utilize the standard FermiMail Services.

2.7.1 Chat

The NOvA collaboration will utilize the standard FermiMail Chat Services.

2.7.2 Email

The NOvA collaboration will utilize the standard FermiMail Email Services.

2.7.3 Mail Lists

The NOvA collaboration will utilize the standard FermiMail Mail List (LISTSERV) Services

2.8 Enterprise Support Services

The NOvA collaboration will utilize the standard Enterprise Support Services including support for an instance of the document management system (DocDB) for use by NOvA.

2.9 Network Services[11]

The NOvA collaboration will utilize the standard Network Services together with the following enhanced services:

- Configuration, monitoring and support of the network switches deployed at the experimental sites, as detailed in Table 1.
- Configuration and monitoring of the WAN connection for Ash River
- Configuration and monitoring of a pseudo-wire between the FNAL campus and the Ash River Laboratory site.
- Configuration, monitoring and support of an integrated, out of band [celluar technology based] connection for the Ash River site.
- Configuration, monitoring and support of wireless access points and associated networks (FGZ and private vlans) at the Ash River and near detector site.
- Configuration of public IPV4 address blocks for the Ash River site with mapping to the fnal.gov domain
- Configuration of private vlans in support of the DAQ, DCS and QA/QC networks at the Ash River and FNAL sites.
- Configuration of access control lists (ACL) for the NOvA DAQ and DCS networks at the Ash River and FNAL sites.
- Wired networking between the Ash River computing center and the following areas
    - Ash River Control Room
    - APD test facility
    - QA workstations deployed in assembly bay area
- Configuration, monitoring and support for DNS, NTP and DHCP services provided locally at the Ash River site.

- Configuration and support for a Power over Ethernet (PoE) switch supporting IP based video cameras at Ash River.
- Configuration of the networking components to support network connectivity of miscellaneous specialized devices deployed at the Ash River site (i.e. Laser Scanning System).
- Networking hardware and hardware components deployed to the Ash River site at "cold" spares.

Support levels for these enhanced services are listed in Table 2.

Table 1 Network switches supported for the NOvA experiment

Responsible	FNAL Service Offering	Monitored Configuration Item(s)
Network Services	Network Facilities Services	r-nova-fd-01, r-nova-fd-02, s-nova-fd-{01…24}
Network Services	Network Facilities Services	w-nova-fd-01, w-nova-fd-02, w-nova-fd-03, w-nova-fd-04,w-nova-fd-05, w-nova-fd-06
Network Services	Network Facilities Services	
Network Services	Network Facilities Services depends on Esnet, U of M	r-s-bdr (r-nova-fd-01, r-nova-fd-02 already included above)
ESNET	N/A	
U of M	N/A	
Networking	Wide Area Network Infrastructure Services	r-nova-fd-03
AT&T	N/A	r-nova-fd-03
Networking		s-exp-nova-1, s-exp-nova-lcc-1, s-nova-lcc-2, s-exp-nova-wh12w

Experiment Component	Service Availability Schedule	Support Availability	Incident Response	Incident Resolution	Request Response	Technical Notes
Ash River/Far Detector Infrastructure LAN	24 x 7	24x7	4 hours 24 x 7	Map to Foundation - High	Map to Foundation	Redundant routers and uplinks
Ash River LAN: End-System Connections	24 x 7	8 to 17 by 5	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
Ash River DAQ LAN: End-System Connections	24 x 7	8 to 17 by 5	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
Ash River Wireless LAN	8 to 17 by 5	8 to 17 by 5 (run pas B.Miller)	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
Ash River VLAN(s)	24 x 7	8 to 17 by 5	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
Ash River to FNAL Primary WAN Connectivity	24 x 7	24x7	4 hours 24 x 7	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
Ash River to FNAL Primary WAN Connectivity			N/A	N/A	Map to Foundation	
Ash River to FNAL Primary WAN Connectivity			N/A	N/A	Map to Foundation	

Ash River to FNAL Backup WAN Connectivity	8 x 5	8 to 17 by 5	Map to Foundation - High	21 days	Map to Foundation	Single point of failure. No local spare.
Ash River to FNAL Backup WAN Connectivity			N/A	N/A	Map to Foundation	
FNAL NOvA Near Detector Infrastructure LAN	24 x 7	24x7	4 hours 24 x 7	Map to Foundation - High	Map to Foundation	Redundant routers and uplinks
FNAL Near Detector LAN: End-System Connections	24 x 7	8 to 17 by 5	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
FNAL Near Detector DAQ: End-System Connections	24 x 7	8 to 17 by 5	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Single points of failure w/local spares
Ash River Primary DNS, NTP servers	24 x 7	24 x 7	4 hours 24 x 7	Map to Foundation - High	Map to Foundation	Fall over to FNAL services
Ash River PoE hardware	24 x 7	8 x 5	Map to Foundation - Low	Map to Foundation - Low	Map to Foundation	
Network ACL Configurations	24 x 7	8 x 5	Map to Foundation - High	Map to Foundation - High	Map to Foundation	Subject to change management. Authorization through NOvA run coordinator.

NOvA is responsible to notifying Fermilab of changes to NOvA's requirements or new computing deployments as early as possible. NOvA should be aware that significant lead-time may be necessary should there be a need to change an existing service or current infrastructure to accommodate NOvA's needs.

2.10 Networked Storage Hosting[12]

The NOvA collaboration will utilize the standard Networked Storage Hosting Services listed below.

2.10.1 NAS

2.10.1.1 BlueArc

The NOvA collaboration uses the following BlueArc volumes:

- Blue3:/nova/data
- Blue3:/nova/prod
- Blue3:/nova/ana
- Blue3:/nusoft/data
- If-nas-0:/nova/app
- If-nas-0:/nusoft/app
- Blue2:/fermigrid-fermiapp
- Blue2:/fermigrid-app

2.10.1.2 (AFS)

The NOvA collaboration uses the standard (/afs/fnal.gov/files/home/*) AFS space for the home directories of its members.

2.11 Service Desk[13]

The NOvA collaboration will utilize the standard Service Desk Services. The Service Desk Service SLA describes the expectations and responsibilities of the customer (NOvA) and the Computing Sector

2.12 Video Conferencing[14]

The NOvA collaboration will utilize the standard Video Conferencing Services.

- Control Room Video conferencing for remote collaborators

3 Scientific Services

3.1 DAQ & Controls

In the areas of *artdaq* and *art*, the experiment will require support for the following:

- *artdaq* utilities libraries: Support will include problem analysis, testing, and bug fixing for only the pieces used for the coupling of the buffer node system and the *art* framework.
- *art* framework: Support for the *art* suite is covered in the offline software section of this document. SCD will provide the same support within the trigger context, since it is the same software that is used. In additional, SCD will provide the same level of support for the integration of the *artdaq* utilities within the *art* framework. Feature requests and requests for changes in requirements will be supported using the standard *art* issue tracking system.

Lists of software components and their disposition are provided in the next several sections. It should be noted that there are no DAQ or controls software components that have been developed outside of SCD that will transition to being supported by SCD.

3.1.1 DAQ components built under the NOvA Project that will be supported by other groups

The areas of responsibility that will transition to groups outside of SCD include the following:

- hardware interface libraries (e.g. the DCM Program Utilities package)
- system diagnostic packages
- software utilities (e.g. NOvADAQUtils and NOvATimingUtils)
- the DCM boot loader
- the message passing interface libraries (RMS)
- management of OpenSplice DDS configurations and any rebuilds of the OpenSplice software libraries that may be needed (taking into account local patches to the software)
- scripts and applications for running the system (DAQOperationsTools)
- the Buffer Node Event Builder application
- maintenance and support of the software that obtains data from the Buffer Node Event Builder and feeds the *artdaq* utilities leading into the *art* framework.
- the Configuration Management libraries, applications and database schema
- the Application Manager application

- the EPICS slow controls subsystem
- customizations and debugging of the SRT build system used for the online software
- management of the NOvA DAQ installation of the Screen Snapshot Service

Unless otherwise noted, the responsibility for the software will include maintenance (e.g. bug fixes and minor enhancements) and support (e.g. configuration changes and the debugging of problems).

During the transition phase, members of SCD will provide documentation, mentoring, consultation, and continued development effort as agreed upon with members of the experiment.

Once the transition phase is complete, members of SCD will be available for consultation on design and implementation details, and available to provide guidance for debugging problems, on an 8x5 basis.

3.1.2 SCD-built components that will continue to be supported by SCD

The areas of responsibility that will remain within SCD include the following:

- DCM/TDU device driver(s)
- DCM/TDU Linux kernel(s) and kernel modules
- the PPC cross compiler
- *artdaq* – the necessary support that is described above (intro to section 3.1)
- *art* – the necessary support that is described above (intro to section 3.1)
- the core Message Facility package

This support will be provided on an 8x5 basis, and any requested work will need to be scheduled based on the priority of the request and the priority of other ongoing work.

3.1.3 Components that will be supported by the experiment

The areas of responsibility that will remain with members of the experiment include the following:

- the package that supplies trigger results from the *art* framework to the NOvA Global Trigger software
- the data driven trigger algorithms
- the data-driven trigger package that integrates the *art* framework, the *artdaq* utilities, and all other DAQ systems for management and monitoring
- all other components that were developed by experimenters or other groups outside of SCD

3.1.4 Ongoing SCD development efforts

In addition to the transition and support activities that will be provided by SCD personnel, there are several areas in which SCD folks will continue to contribute development and deployment effort as the experiment commissions the DAQ and slow controls systems for the Near and Far Detectors. These include:

- setup of the dynamic hostname resolution scheme for DCMs at the Near Detector, if needed
- development, deployment, and debugging of the DCM firmware and software that is needed to support multi-point readout
- deployment and integration of the EPICS-based slow controls subsystem at the Near Detector
- any enhancements to SCD-developed applications that may be needed to scale to the large number of channels in the full Far Detector

These efforts will proceed on an 8x5 basis, similar to the construction effort that has taken place so far.

Additional requests may arise as the commissioning of the detector and DAQ continue. The breakdown of responsibilities, and schedule for addressing each request, will be negotiated on a case-by-case basis.

3.1.5 Computing Sector Responsibilities

SCD support for DAQ and controls software in year 1 will include the effort required to bring the system to full operation. This is expected to be a large commissioning effort. Beyond the first year the support is limited to that required for baseline operations. Consulting services for major changes or upgrades can be provided upon request by NOvA and agreement with the relevant departments.

- Provide consulting services to NOvA online support personnel regarding best practices, technical issues, or specific problems related to system administration of online computing systems. CS will provide limited technical assistance to deal with major system administration problems. All such services will be provided on an 8x5 basis;
- Provide ongoing development and support effort as outlined above.
- Provide system administration and support for DAQ/online systems in a manner consistent with the 24x7 nature of data taking and the high value of beam time and data taking. The potential for transitioning some of this support to SCD in the long-term will need to be negotiated with CS.
- DAQ operations and computing support for facilities located at the Far Detector in Ash River, MN.

3.1.6 NOvA Responsibilities

- Continuing to maintain and support the software components that were developed outside of SCD.
- Assume responsibility of the maintenance and support of the software components that will no longer be the responsibility of SCD, as outlined above.
- Comply with security requirements outlined in the NOvA Minor Application Plan;

3.1.7 Joint Responsibilities

- Communicate feature requests, problems, and the availability of new versions of core packages in a timely way.

3.2 Engineering and Electronics

The NOvA collaboration will utilize the standard Engineering and Electronics Services as well as the following enhanced services:

SCD support for DAQ hardware and firmware in year 1 will include the effort required to bring the system to full operation. This is expected to be a large commissioning effort. Beyond the first year the support is limited to that required for baseline operations. Consulting services for major changes or upgrades can be provided upon request by NOvA and agreement with the relevant departments.

1. Provide electrical design and firmware revisions for the DCM, MTU, and TDU required to maintain operation of the full DAQ chain within design specifications;
2. Take on responsibility for the FEB, HV/LV power supplies and the PDM units.
3. Provide a swap and repair service for the modules in 2. and 3. Above. This will be provided via the Logistics counter staffed by Dell Managed Services.
4. Provide hardware support for DAQ network switch arrays.

The DAQ for NOvA includes several electronics modules that need custom support and sparing. These are listed explicitly below:

3.2.1 Custom Front End Electronics:

- Fermilab is discussing support for the Custom Front End Boards developed by Harvard
    - Fermilab would require a test-stand, drawings, documentation for testing the modules
    - A complete count of spares and location of those spares is essential
    - Fermilab ESE department would provide hardware support & repair (fix when broken)
    - The logistics of swapping out broken FEB's and replacing them with good FEB's would be provided by the Fermilab Logistics Counter through Dell Managed Services
    - Firmware support (bugfixes and feature set changes) would be discussed and taken on if appropriate in the future.

3.2.2 Data Concentrator Modules (DCMs):

- Fermilab ESE department provides hardware support & repair
- Logistics for swapping out broken DCM's are similar to FEB's
- Firmware support (bugfixes and feature set changes), as agreed
- Configuration and design support for Linux kernel drivers for the device are provided

3.2.3 Timing Distribution Units (TDUs):

- Custom modules developed by Fermilab
- ESE hardware support (i.e. fix when broken)

- Firmware support (bugfixes and feature set changes)
- Custom Linux kernel drivers for the devices (bugfix)
- Custom ARM processor software (bugfix and feature set changes)
- GPS antennas and receivers (commercial units, route to vendor for repair/replacement)
- Timing check hardware (pps boards)

3.2.4 Power Distribution Modules:

- Custom modules developed by the University of Virginia
- Fermilab Logistics hardware support & repair (fix when broken)
- Fermilab ESE proposes to take on the repair of the modules, will require drawings, documentation, inventory of spares

3.3 Offline Computing

The NOvA collaboration depends on the standard Grid and Cloud Computing Services. Scientific Computing System, Scientific Data Storage and Access, and Grid and Cloud Services provide support for NOvA data analysis and processing systems under the Foundation SLA with 8x5 support. The number of batch slots, experiment data storage size and performance, and common job submission and monitoring tools are provided as part of these services. The needs for each year are proposed, agreed to, and captured through the Fermilab Scientific Portfolio Management process. This information is recorded at **http://tinyurl.com/SPPM2014**.

3.3.1 FermiGrid

The NOvA collaboration uses the standard FermiGrid Services. The NOvA experiment relies on FermiGrid as the ensemble of interfaces and services to access the Fermilab computing infrastructure. The experiment is in the process of targeting different computing platforms for different computing tasks. The data-intensive computing activities, such as reconstruction, target mostly FermiGrid; purely compute-intensive tasks, such as monte-carlo production, target mostly distributed resources, such as OSG or public and commercial clouds. CS takes responsibility to manage the ensemble of the services that allow access to the computing infrastructure at Fermilab at the level described in the SLA.

3.3.2 FermiCloud

The NOvA collaboration is discussing with the Computing Sector use of Virtualized and Cloud Services.

3.3.3 GridFTP

The NOvA collaboration uses GridFTP Services. For transferring the output of Grid jobs, NOvA relies on a Globus GridFTP server configured to maintain both user and *group* id file ownership. The group id ownership is particularly relevant for this service because other data trnsfer services do not necessarily preserve it.

3.3.4 Accounting Service

The NOvA collaboration will use the standard Gratia Accounting Services.

3.3.5 Jobsub

The NOvA collaboration will utilize the standard FIFE Jobsub Services. JobSub is an ensemble of services to submit and manage jobs to local and remote resources. The ensemble includes a user-facing interface for job management, which encapsulates the semantic of experiment-specific use cases, job queuing and resource matching services, basic provisioning services, as well as input / output sandbox transfer service. NOvA relies on this service for the submission of all jobs to resources either local or remote, dedicated or opportunistic, public or private or commercial.

3.3.6 Fifemon

The NOvA collaboration will utilize the standard FIFEmon Services. Fifemon is the service that monitors the status of submitted jobs. In contrast to the Gratia accounting service, which shows the resources used by competed jobs, Fifemon shows the status of the jobs as they go through their lifecycle e.g. submitted, idle, running, completed. The service allows to "drill down" at an increasing level of detail for those jobs of particular interest. As NOvA undertakes increasingly complex computational campaigns, it will rely on this service more and more.

3.3.6.1 Computing Sector responsibilities

1. Operation and support for use of the Interactive/Batch Analysis Cluster GPCF.
2. Operation and support for use of local Grid accessible resources agreed to with the experiment.
3. Support and consulting for the use of offsite resources through the Open Science Grid.
4. Provide consultation with offline personnel from the experiment on issues related to grid utilization.
5. Develop and provide training and documentation in the recommended use patterns of the above resources.

3.3.6.2 NOvA responsibilities

6. Validate users authorized to access NOvA grid computing resources. The experiment will further provide personnel for the roles of "Group Managers", "Operations Contact", "Security Contact" and "Spokesperson", pursuant to the "Establishing Grid Trust with Fermilab" document [3].
7. Document the local grid and interactive CPU resources required to meet the physics goals of the experiment.
8. Ensure that NOvA users are informed as to the appropriate usage patterns for all CPU resources[1]. Work with CS personnel as needed to investigate and address operational issues or utilization efficiency problems.

[1] Experiments that use Grid resources must establish the appropriate Grid Trust Agreements [3] prior to use of the Fermilab Campus Grid (FermiGrid) resources. In addition to the Fermilab Policy on Computing, specific additional policies apply to Grid computing activities on FermiGrid [4] and further policies apply to Grid resources accessed via the Open Science Grid (OSG) collaboration [5].

9. Perform job submission and data processing tasks.
10. Provide user support for job submission and job tracking, and user documentation and education on the use of NOvA computing resources.
11. Provide those components of a job submission layer to the batch and grid resources that is specific to NOvA.
12. Specify and develop any monitoring capabilities that are needed to effectively utilize CPU resources, but that are not provided by available monitoring tools. Instrumentation of NOvA executables or glide-ins are possible examples where joint effort may be required.
13. Provide feedback on the training and documentation provided by the Computing Sector.

3.3.6.3 Joint responsibilities

14. Meet as needed to discuss operational issues affecting the use of computing systems, best practices for using the systems, user support issues, utilization strategies, or other items of mutual interest with respect to the computing systems.
15. Investigate and deploy suitable mechanisms for transferring executables, database information, etc., to remote worker nodes for the purpose of Monte Carlo generation, and for transferring generated files back to Fermilab.

3.4 PREP

The NOvA collaboration will utilize the standard PREP Services

3.4.1 Prep Logistics

The NOvA collaboration will utilize the standard PREP Logistics Services. PREP support is through standard replace and repair procedures with the availability of the Prep service window being 9.30am -4pm 5 days a week. All PREP loans are authorized under a TSW. Experiments sign full TSW's with the division heads organized by the Directorate Office of Program Planning. Test beam experiments do the same, save that the CS signature has been delegated by the Division head to the PREP Scientific Manager (PSC).

3.4.2 PREP Electronics

The NOvA collaboration will utilize the standard PREP Electronics Services, together with the following list of additional enhanced services:

There is a TSW template for offsite loans signed by the User, PREP Scientific Manager, and Associate Director for Program Planning. Expansions beyond the "PREP list" in an TSW are normal, expected, and by negotiation. There are no explicit Service Level Agreements (SLA's). Implicit in the pool model is that working spares are available to replace failures and diagnose problems. PREP, when asked, will do whatever it can to get a running experiment, which is down back to taking data. This includes spares, replacements, and technical consulting with the Techs and managers as required.

The NOvA experiment will utilize a broad array of electronics in the development, commissioning, and operation of the DAQ system. The equipment required includes standard test and laboratory equipment (e.g., oscilloscopes, voltage meters, current load boxes, NIM crates and associated modules), basic data acquisition systems needed to

interface with other laboratory systems (e.g., CAMAC crates and modules needed to receive and process signals from the accelerator facilities), and NOvA-specific hardware procured from outside vendors or built in-house. Some details of these special cases are described below.

3.4.2.1 Prep Responsibilities

6. Low Voltage Power supplies: The Near Detector requires 4 operational supplies plus 3 spares. The spares may be used by Prep for testing. The Far Detector will employ 66 operational supplies and have 4 spares.
7. High Voltage Power supplies: For the near detector the number of supplies that will be procured is 3, where 2 of them are denoted as spares. The NOvA experiment will use floating high voltage power supplies from WIENER Plein and Baus Ltd. The Near and Far Detectors each use 1 supply and have 2 spares.
8. Data Concentrator Modules: For NDOS, there are 11 modules installed. For the Near Detector there will be 14 modules and 4 spares. For the Far Detector there will be 168 modules and 20 spares.
9. Timing systems: NOvA uses two timing devices, the Master Timing Unit (MTU) and the Timing Distribution Unit (TDU), to provide synchronization of the front end boards. For the Near Detector there will be 2 masters and 4 slaves. For the Far Detector there will be 2 masters and 28 slave units. Provide and maintain standard test and laboratory equipment as agreed upon with the experiment.

3.4.2.2 High Voltage and Low Voltage Power Supplies:

- Commercial modules
- Fermilab Logistics for swapping bad units and route to vendor for repair/replacement as necessary.

3.5 Scientific Collaboration Tools

The NOvA collaboration will utilize the standard Scientific Collaboration Tools Services.

3.5.1 Redmine

The NOvA collaboration depends on the standard Redmine Services.

3.5.2 CVS/Subversion/GIT

The NOvA collaboration will utilize the standard CVS/Subversion/GIT Services. NOvA code repositories are hosted through cdcvs.fnal.gov redmine core repository and collaboration management system.

3.5.3 ECL

The NOvA collaboration will utilize the standard ECL Services together with the following enhanced services.

The NOvA collaboration depends on the electronic log-book application for both the near and far-detector. Support for the database servers Support for the latter is needed round the clock during data taking. Tickets for off-hours support for ECL database servers will be generated by the PPD Experimental Operations support organization (IFTBG).[15]

See PPD support documentation for enhanced support levels.

3.5.4 UPS/UPD

The NOvA collaboration will utilize the standard UPS/UPD Services

3.6 Scientific Computing Systems

The NOvA collaboration will utilize the standard Scientific Computing Services. These services include support for the NOvA collaboration DAQ systems (OS installation and patching) and other NOvA collaboration desktops.

3.6.1 Experiment Desktops and Control Room Workstations

The NOvA collaboration will utilize the standard Experiment Desktop Services. The NOvA experiment utilizes a set of 10 workstation class desktop computers that are used in the NOvA control rooms for control both at Ash River and in the WH12 control room for control and monitoring of the experiment. These machines are designed and configured to be generic display stations (i.e. any machine can display any DAQ desktop or interface) and no single machine is considered a critical system for operations. These control room workstations are supported under the SLA for Nova control room computing.

3.6.2 DAQ Computing Clusters

The Far Detector DAQ computing hardware is composed of a cluster of commodity computer systems. The Far Detector DAQ computing is additionally configured with infrastructure to provide each system with:

- Serial console port access,
- Remote power on/off via network controllable PDUs,
- Full access (to the bios level) via keyboard/video/mouse servers that are accessible via TCP/IP.

3.6.2.1 Near Detector Computing Systems:

- Hardware support (replacement under warrantee) for approximately 25 worker nodes
- Hardware support (replacement under warrantee) for 2 disk arrays
- Hardware support (replacement under warrantee) for support infrastructure (consoles, PDUs, etc.)
- System administration for approximately 27 SLF6 systems, including configuration management, software updates and security Patches

3.6.2.2 Far Detector Computing Systems:

- Hardware support (replacement under warrantee) for approximately 200 worker nodes
- Hardware support (replacement under warrantee) for between 3 and 6 disk arrays
- Hardware support (replacement under warrantee) for support infrastructure (console servers, PDUs, etc.)
- System administration for approximately 210 SLF6 systems, including configuration management, software updates and security patches

3.6.2.3 SLA and deviations

1.The Full SLA for scientific servers can be found here.

2.The computer security details are provided in the Minor Application Plan Computing Sector responsibilities

- Installation of, updates, security and other patches for the Scientific Linux OS
- Monitoring and system administration services
- Installation and support of the PUPPET configuration management software.
- hardware is under maintenance contract with the corresponding equipment's vendor, and any hardware located at the Far Detector that is found to require repair will be shipped from Ash River to Fermilab for warranty maintenance.

3.6.2.4 NOvA responsibilities

- Install and support of all online application software, and Fermilab supported physics toolkits and utilities needed.
- Provide schedules for deploying security patches to all systems that are consistent with Lab security policies and the NOvA Minor Application Plan.
- Provide an expert either from from the experiment or from PPD oncall support who can assist system administrators.

3.6.2.5 Joint responsibilities

Any system or support level Change planning, requests and documentation

3.6.3 CVMFS

The NOvA collaboration will utilize the standard CVMFS Services.

3.6.4 Build Service

The NOvA collaboration will utilize the standard Build Service Services (once it is commissioned).

3.7 Scientific Data Management

NOTE: Scientific Data Management will be on-boarded to ITIL later in FY14, and so currently there is no Service Level Agreement (SLA) in place. This TSW will serve in lieu of the SLA. Once the SDM SLA is accepted, that will supersede the statements in this section and in subsections below.

The Scientific Data Management services involve management of the experiment's event data files and includes the following service offerings:

- SAM/IFDH
- File Transfer Service (FTS)
- Data handling

The guiding principle is to provide data handling and management services to an experiment that are robust, efficient, innovative, easy to use, and easy to maintain and operate, and low cost.

The suite of offerings enables the NOvA experiment to catalog and store event files in the Fermilab central storage system and retrieve such files for processing by jobs running at Fermilab and at remote sites.

Note that the term "event data files" is used here to describe the type of files handled by these services. These files generally contain event information originated by the NOvA detectors or simulation. They may be Root files or based on a private format. Individual log files, histogram files, documents, and such are generally not handled by SDM services. There are exceptions listed here,

- *art* configuration (FHICL) files for simulations are handled by SDM services. Such files are used to initiate simulation jobs and serve as the top "ancestor" of files produced by the simulation run

- Log file bundles (e.g. tarred and compressed) may be handled by SDM services for archival purposes. Such bundles should be large (>2 GB) if possible.

  Other file types may handled by SDM services upon mutual agreement by NOvA offline management and the SDM services management and should be listed above.

SDM services are generally geared for access to data from batch jobs. There are situations described below where interactive access to data may be possible.

3.7.1 SAM/IFDH

SAM is a system that catalogs, moves, and tracks event data from a central storage facility to caches world-wide. IFDH is a complementary system that handles the "last mile" of data transfer from a nearby cache to the worker node hosting the running job. SAM provides an interface (SAMWeb) for the user and administrators to configure, communicate with, and monitor the data management system.

The current configuration of NOvA's data access system is to use Fermilab's Enstore tape system for archival storage and the central cache system with dCache for cache management. The SAM system involves software as well as several physical services including SAM stations, stagers, and other servers. SCD will maintain such software as well as operate the physical services. There are two types of incidents that may be raised against the SDM services: regular incidents are those with either low impact or low urgency (as defined in the CS Foundation SLA). Critical incidents are those with both high impact and high urgency and must be initiated by the NOvA offline manager or delegate and called into the Service Desk. Examples of incidents are

- an outage of a critical SAM physical component that halts data handling for all jobs at all sites is a critical incident
- an outage of a web service that prohibits access to all of data management for the experiment is a critical incident
- a software bug that causes a race condition halting all data handling for all jobs at all sites is a critical incident
- because the FermiGrid site is the primary batch system for NOvA processing, a problem that halts data handling for all jobs on FermiGrid is a critical incident
- a remote site problem that causes data handling to fail at that particular site is a regular incident
- a problem that halts interactive access to files but leaves jobs unaffected is a regular incident

- problems of low impact (affecting only a few people) or low urgency are regular incidents
- in special circumstances, a regular incident may be elevated to critical with mutual agreement between the NOvA offline manager and SDM service managers.

As mentioned above, SAM and IFDH are geared for data delivery to batch jobs. For some situations, IFDH will have capabilities to deliver files to interactive sessions. Currently, these situations include:

- Interactive sessions on Fermilab NOvA GPCF interactive nodes (e.g. novagpvmXX) with the file resident or accessible to the central public dCache.

Data transfer rates and requirements are specified in SPPM proceedings.

SCD Responsibilities:

- Provide the Scientific Data Management services that enable NOvA to catalog, store, and retrieve event data as described in this section.
- Support SAM and IFDH software, interfaces, and libraries at a 8x5 level except in the case of bug that causes a critical incident. Such critical incidents are handled at an 8x7 level.
- Support SAM physical services and servers at an 8x5 level except in the case of an outage that causes a critical incident. Such critical incidents are handled at an 8x7 level.
- If a remote site experiences no or under-performing file delivery, the site will be investigated and debugged by SCD. Note that cooperation with site administrators may be necessary and NOvA managers may need to assist.
- Respond in a timely manner to requests generated by NOvA.
- Monitor the system for operations and performance at an 8x5 level. As practical as possible, critical problems generate an automated page to SDM service operations personnel. Note that not every problem can be anticipated nor covered by automation. In the case of a critical problem page due to or resulting in a critical service outage, the SDM service operator will notify the NOvA offline manager or delegate of the outage.
- Perform maintenance on systems and software as necessary. Such maintenance may incur a service outage or degradation during the maintenance window. Such maintenance windows must be negotiated with NOvA offline management in advance. SDM service management will make every attempt to minimize the occurrence of unplanned emergency maintenance windows.

NOvA Responsibilities:

- Open Service Desk incidents when problems are noticed. Critical incidents as described above need to be initiated by the NOvA offline manager and called into the Service Desk.
- Open Service Desk Requests for new meta data fields or advice about meta data fields.
- Open Service Desk Requests for new use cases or anticipated unusual increase in demand as early as possible. The SDM service management and operators will respond to the request and if possible, adjust the services configuration accordingly.

- While the SDM systems will be as robust as practical, NOvA offline management should prevent, as much as practical, user abuse of SDM systems that cause unwanted increases in demand or unusual use cases.
- Negotiate with SDM service management for maintenance windows. Note that unplanned emergency maintenance windows may be necessary in special circumstances.

3.7.2 File Transfer Service (FTS)

The File Transfer Service (FTS) is a robust system for uploading files into the SAM catalog and central Fermilab storage. Its main use is for transferring event files from the NOvA DAQ systems into SAM as well as the output of Simulations. The FTS uses many SAM components and thus SCD and NOvA responsibilities detailed in the SAM/IFDH section apply here. FTS also introduces its own software and physical server components and those are supported at a level similar to the SAM software and services. Both SDM service management and NOvA operations will monitor the FTS system through its built in web monitoring service.

Since files produced by the detector DAQ systems are irreplaceable, NOvA will provide enough "spool" disk area to allow for the storing of at least a week's worth of data in case of a full and long term outage of FTS, SAM, or Enstore (such an outage is not anticipated, of course).

NOvA will create Service Desk incidents if problems are noticed with FTS. Critical incidents may be initiated by the NOvA online or offline management and called into the Service Desk. An incident that stops all storage of data (from DAQ or simulation jobs) is critical and will be handled at an 8x7 level.

3.7.3 Data Handling

Data handling is a service offering where SDM experts provide consultation and advice for an experiment's data handing needs aligned with the principle stated above. Examples of topics include:

- Definition of file meta-data
- Exploring new data handling use cases and paradigms
- Exploring data handling technology

NovA may initiate consultation by opening a Service Desk request. SDM experts may approach NOvA offline management for discussions as well.

3.8 Scientific Data Storage and Access[16]

The NOvA collaboration will utilize the standard Scientific Data Storage and Access Services. The Scientific Data Storage and Access services are described in the related SLA. Support for NOvA falls within the standard service categories. The responsibilities of each of the parties is described in the SLA. The expected scale and performance of the systems is described in the submission to the Laboratories Scientific Portfolio Management process.

The disk and tape storage needs of the NOvA experiment are categorized below.

3.8.1 Raw Data

The NOvA experiment records data related to the operation of the near and far detector corresponding to the NuMI beam spills as well as additional data for calibration and for

non-accelerator analysis topics. The data rates for NOvA detector depend on the operational parameters of the detectors (noise floors and zero suppression thresholds) but have been measured using the Far Detector as well as calculated based on future scaling of the experiment. With the current detector configurations the NOvA detectors, scaled to the full 14-kiloton detector size, are estimated to produce a minimum of 8 TB/year of NuMI beam trigger data along with an additional 0.4 PB/year of calibration data. The maximum amount of data that the NOvA detectors can produce per year is estimated to be 43 TB/year for the NuMI beam trigger data and an additional 2.2 PB/year for calibration and other off-spill triggers. These numbers are based on studies of the bandwidth limits of the DAQ system.

The NOvA far detector site is able to locally store/buffer 30 days of detector data on disk arrays purchased by the NOvA project and housed in the Ash River computing center. Under normal operations these disk arrays serve as a temporary storage location for the data during the time when it is being transferred to FNAL for archiving and data processing.

Under the current NOvA computing model, raw data from the detector is transferred using the FNAL developed "File Transfer Service" (FTS) between the Ash River site and a set of data handling stations (novasamgpvm01-0X) located at Fermilab. The data is stored temporarily on the BlueArc central disk system, where it is checked for transmission errors and cataloged by the SAM data catalog system. The data is then copied to the enStore tape system for archival storage as well as to the central dCache pools for general processing and analysis access. For raw data, two copies are made of the data (on different physical tapes) to ensure against data loss in the advent of media failure. The buffered copies of the data are removed from the BlueArc and Ash River disk arrays after the transfers to tape are verified.

Under this raw data model, when the detector enters full production, a raw data volume of at least .8 PB/year will be written to the FNAL tape systems. In addition the system will require an allocation of dCache write pool sufficient to handle the writing of the raw data stream to the enStore system. This size of the pool is estimated at 30-50TB and is sized to accommodate both the nominal [steady state] transfer of data from the far detector, as well as the non-standard flow of data that would be required after a prolonged network outage of the far site (i.e. catch up processing). The system currently also requires 30-50 TB of central BlueArc storage to serve as a buffer for doing the data integrity checks which are performed by the FTS. This storage can be retired when the FTS is upgraded and integrated directly with the dCache pools.

3.8.2 Monte Carlo Simulation Files

The NOvA experiment needs to generate significant amounts of simulation to perform different analysis. This simulation is broken down into different Monte Carlo sets depending on the analysis being performed. The Monte Carlo simulations are generated through the use of a neutrino interaction generations package combined with a GEANT4 based detector simulation package. The output of the initial Monte Carlo generation/simulation step is saved to disk/tape for further processing. For the initial FY14 first analyses the amount of simulation that is required is estimated at 242 TB of tape with an active disk footprint of 70 TB. This simulation set corresponds to 2.5E7 event triggers that are equivalent to 2.5E24 Protons on Target (POT).

3.8.3 Production Data Analysis

The production data analysis is expected to inflate the data volume by a factor of 2-4 over the initial raw or Monte Carlo files sizes. This increase is already included in the estimates

for the required space for simulation. For the initial analysis, the full production chain will only be applied to the beam data and to 1/10th of the off-spill data. The estimate for the full size of this data is 54 TB/year.

3.8.4 End User Data Analysis

The data used for final, high level, analysis by the NOvA collaboration will take the form of small specialized "skims" of the full production data sets. These datasets are expected to achieve a data reduction ratio of at least 100:1 for the primary physics analysis. The storage requirements are estimated to grow from less than 1 TB up to approximately 10 TB total over the course of the first few years of data taking. The skim data volume when combined over all analysis groups and efforts are expected to be less than 100 TB. The expectation is that these skims will reside initially on central disk systems that allow for some degree of "interactive" access to the data (i.e. accessible in a ROOT analysis session.) as well as be accessible from grid/cluster computing.

Computing Sector Responsibilities

1. Install and maintain a central disk pool capable of serving NOvA data to the the GP Grid Farm, CPCF cluster, and other on-site NOvA computers via NFS. Data serving rates must be sufficient to meet the demands of reconstruction and analysis on the GP Grid Farm, GPCF, and other on-site computers. Estimates for this scale are documented elsewhere. The data volume starting in 2014 is expected to be about 30 TB, with subsequent years evolving as indicated in Table 1.
2. Install and maintain NFS-mounted disk serving software releases to the GPCF cluster, GP Grid Farm, and other on-site interactive machines, and machines with disk for building software releases. The experiment expects that a software release area of 2 TB will be adequate through the Near Detector ramp-up phase.
3. Install and maintain project disk to support analysis activity. At present, this disk is provided as part of the GPCF plan.
4. Provide a tape data archive via Enstore accessible. All raw detector data, processed data, and Monte Carlo data will be archived to tape. The FY2010 volume is expected to be about 100 TB, as shown in Table 1. Typical daily storage rates are expected to vary between 100 GB and 1 TB.
5. Provide AFS-mounted home area disk. Regular backups of this space will be performed. Currently, the experiment uses about 100 GB of AFS space for home areas;
6. Except as specified below, monitor performance of tape and disk storage systems.
7. Provide tools for archiving analysis data.

3.9 Scientific Databases

The NOvA collaboration will utilize the standard Scientific Database Services. Scientific Databases and Database Hosting Services are both relevant to the support levels for the experiment systems. Unless otherwise stated the support level is 8x5 as stated in the Fermilab Foundation SLA.

The NOvA experiment employs a diverse set of databases that are used to store information regarding both the construction of the experiment and the operations of the detectors and the conditions of beams the detectors see. These databases are used by both online and offline systems with different access patterns, replication needs and uptime requirements. Details of all databases used only by the NOvA Project and Experiment is found in NOvA-doc-10602. The relevant information for this documents are characterized in the sections below.

The major database applications include the data management catalog, the conditions database, construction and hardware database, and experiment logbook. In the following summary of responsibilities, support for a database application by an organization implies support at all three levels unless specified otherwise.

In addition to the database resources located at Fermilab, database servers are installed at the Far Detector site. The requirements for supporting the installation, commissioning, and operation of the Near Detector will be similar to those for the Far Detector. Resources will be provisioned accordingly as Far Detector installation begins.

CS provides tiered levels of support for database services ranging from 8x5 to 24x7, the choice depending upon the application and the type of underlying database. For support under the Database Hosting SLA, it is required that each production database instance will be accompanied by corresponding integration and development instances, both of which receive the lowest available tier of support.

3.9.1 Database Descriptions

The database applications that are required for the construction phase of the NOvA experiment are the UMN module factory database, Ash River construction database, and FNAL hardware database. Each of these PostgreSQL databases contains quality assurance (QA)/quality control (QC) and installation data on all detector components, and are critical to the construction project. Data are replicated between these three databases by the experiment. NOvA requires 8x5 support for only the FNAL hardware database and the corresponding web interface application until at least the completion of the construction of the NOvA Far and Near detectors (mid-2014).

The web-based "Conditions Database" service is an application that is also used by the NOvA offline. The Conditions Database web server and application require 24x7 support.

NOvA also uses the IFBEAM and the SAM data management databases. The IFBEAM database is used by the NOvA experiment in both real-time and in an offline capacity. The IFBEAM monitoring web application is used in real-time and is important for detector operations. Both of these databases and web services require 24x7 support.

The general operational parameters for these sets of databases are summarized in the table below:

Table 2 - Nova Database Properties

Database	Type	Access Types	Support Requirement
Hardware	Postgres	CS Web App	8x5

SAM	??	CS Custom App	24x7
IFBEAM	Postgres	CS Web App	24x7

3.9.2 Responsibilities

3.9.2.1 Computing Sector responsibilities

Install and maintain databases, database servers and applications needed to store and utilize the following mission critical data:

- Support the schema for the NOvA Hardware Database, a web-based GUI to access the Hardware Database, and a set of import/export tools for data in the Hardware Database;
- Install and maintain an instance of the NOvA conditions database web server.
- Install and maintain an instance of the Control Room Logbook as needed by the experiment.

3.9.2.2 NOvA responsibilities

- Enter the content of all databases;
- Interfacing NOvA software with the database applications;
- Ensuring that users are informed as to appropriate usage patterns, and otherwise assisting CS personnel in investigating and addressing operational issues.
- For cases in which there is no existing schema or database application, specify and document the requirements, the use cases and queries needed, etc., as requested by the CS.
- Provide time windows during which regular database maintenance may be performed and security patches applied in a manner consistent with Fermilab security policies and the NOvA Minor Application Plan.

3.9.2.3 Joint responsibilities

- Developing and approving the specifications for user access, the database applications and schemas.
- Participate in annual "Taking Stock" meetings to long-term operational issues and resource planning. CS will coordinate these meetings.

3.10 Scientific Frameworks

The NOvA collaboration depends on the art software suite for their offline production (simulation and reconstruction) applications. The entire suite appears as external packages to NOvA. A portion of the developed applications is also deployed within the Data Driver Trigger (DDT). The suite can roughly be broken up into three areas: the art framework, support libraries, and external products.

The NOvA experiment requires support for the following packages:

- Art suite releases: CS will create, host, and maintain release distribution of the entire art suite. Distributions may include bug fixes, features, or changes build parameters as required by the NOvA experiment release manager. CS will provide support through the standard support systems for integration issues and questions.

Platform support will be primary SLF6. SLF5 will be supported as a secondary platform at low priority.

- Art suite software: CS will provide support for the art framework and underlying support software libraries. Support will include bug fixing, problem analysis, answering questions concerning functionality and usage, and providing upgrades to accommodate platform and external product changes. CS will accept feature requests through the redmine system.
- UPS: CS will provide support for the UPS package. This includes answering questions, investigating problems, and providing bug fixes when problems are encountered.

3.10.1 Supported by the experiment

The Software Release Tools (SRT) is a package is used to build software releases from the NOvA repositories. It is used in the compiling and the running of the offline and online experiment code developed by the experiment.

The experiment shall provide necessary support for the following packages:

- SRT: This includes troubleshooting runtime errors caused by mismatched packages, corrections to compilation rules, and maintenance of the tools and scripts within SRT.

Experiment offline release builds and distributions, including integration with external products.

3.11 Scientific Software

The NOvA collaboration depends on ROOT and TotalView.

3.12 Simulation Software

The NOvA collaboration relies on Geant4 and Genie releases and bug fixes for detector simulation.

4 Miscellaneous

This section is used for miscellaneous items not covered above.

This TSW replaces the following existing documents.

4768-v1, NOvA Construction Project Computing MOU

3597-v7, NOvA MoU with the Computing Division (DEPRECATED)

3528-v1, NOvA MOU with Computing Division (DEPRECATED)

5 References

[1] FNAL Foundation Service Level Agreement, CS-DocDB #4042
http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4042

[2] Fermilab Policy on Computing http://security.fnal.gov/policies/cpolicy.html

[3] "Establishing Grid Trust with Fermilab", CS-DocDB #3429 http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4042

[4] The Fermilab Campus Grid (FermiGrid) Computing Policy Page http://fermigrid.fnal.gov/policy.html

[5] The Open Science Grid http://www.opensciencegrid.org/

[6] "Technical Design Report for CD-2/3a," NOVA Document 2678-v8, October 8, 2007.

[7] NOvA Minor Application Plan (in progress), NOvA DocDB #4250.

[8] The following NovA documents describe the DAQ system requirements: DAQ Monitor (#3769), Data Concentrator Module (Software) (#3664), Data Logger (#3683), Data Quality Monitoring (#3799), Dispatcher (#3944), Event Builder (#1168), File Transfer System (#3786), Global Trigger (#2631), Spill Server (#4529), Message Logging System (#2332), Message Passing System (#1210), Resource Manager (#3678), Run Control (#1877). These documents are in the NOVA-DocDB.

6 Document Revision History

Date	Version	Author(s)	Comments
13-Sep-2013	V0.00	Keith Chadwick	Template
30-Sep-2013	V0.1	Keith Chadwick	Import NOvA information from prior MOU
26-Sep-2013	V0.2	Andrew J. Norman	Edited to reflect Nova Computing
2-Oct-2013	V0.3	Andrew J. Norman	More edits
4-Oct-2013	V0.4, V0.5, V0.6, V0.7	Keith Chadwick, Ruth Pordes	Further edits and refinements, Add replacement text from SSA
30-Oct-2013	V0.8	Ruth Pordes	Organize to follow Service Areas & Offerings Structure throughout
12-Dec-2013	V.9	Margaret Votava	Shuffled things around. Added some missing offerings. And made offerings an outline level underneath services areas
2-Jan-2014	V0.19	Andrew J. Norman	Revisions from review by NOvA
23-Jan-2014	V0.22	Andrew J. Norman	Cleaned up formatting and added database hosting section

7 Bibliography

[1] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4042
[2] Fermilab Policy on Computing http://security.fnal.gov/policies/cpolicy.html
[3] *NuMI Off-Axis nu_e Appearance Experiment Technical Design Report (TDR)* [FERMILAB-DESIGN-2007-01].
[4] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4314
[5] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4315
[6] http://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=4321
[7] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4773

[8] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4664

[9] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=3716
[10] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4310
[11] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4312
[12] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4311
[13] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4591
[14] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4313
[15] PPD support documentation to be completed
[16] http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5032